

DAQ Consortium Update

Dave Newbold for the DAQ consortium
LBNC Review, 19-Feb-18

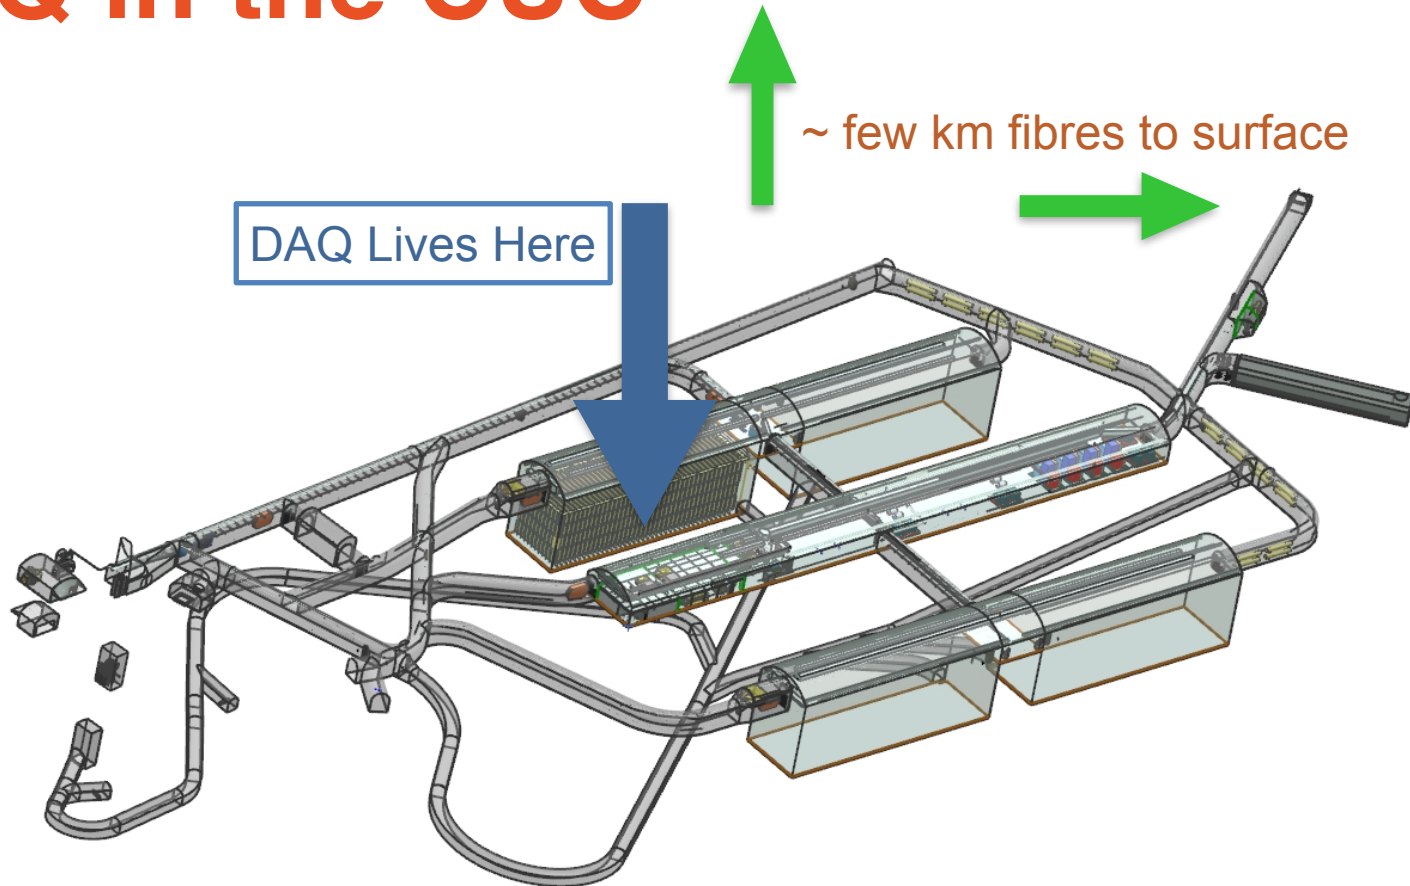
DAQ workshop, Oxford, Jan 2018



DAQ Scope

- Basic functions
 - ▶ Synchronise and **receive** digital data from all FD sub-detectors
 - ▶ **Buffer** all data pending a trigger decision – without information loss
 - ▶ Form ‘**trigger primitives**’, summarising localised observed activity
 - ▶ Make local, then regional, then global trigger **decisions**
 - ▶ **Select** space & time ranges of data, relay them to permanent **storage**
 - Meaning: across the WAN to FNAL and other off-site computing centres
 - ▶ Option: carry out further data **selection** / refinement before transfer
- What’s different and interesting about DUNE?
 - ▶ System has to be extremely **flexible** – ‘permanent commissioning’ mode
 - ▶ There are **special requirements** from the physics – e.g. SNB data buffer
 - ▶ **Hard-to-access location**: reliability and full remote operation are essential
 - ▶ Relatively short **time scale** to design and build an (up to) 50Tb/s system

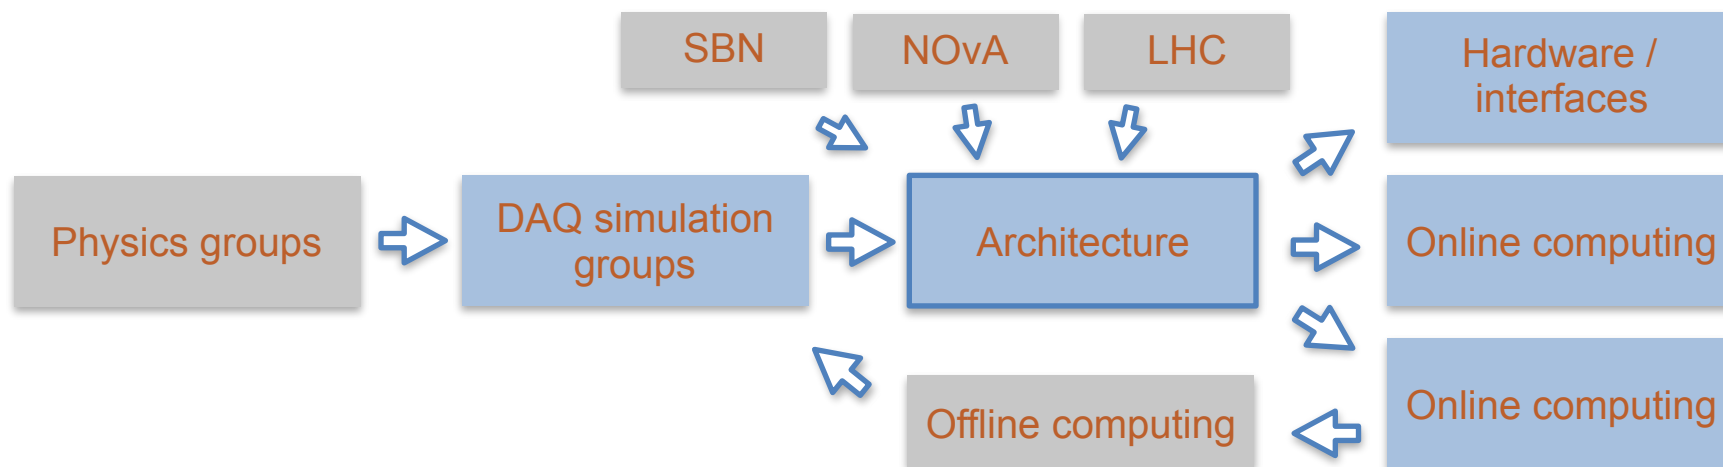
DAQ in the CUC



- ... eventually, after design / prototype / construction / deployment
- But not *just* here – also surface elements, remote operations

DAQ Consortium

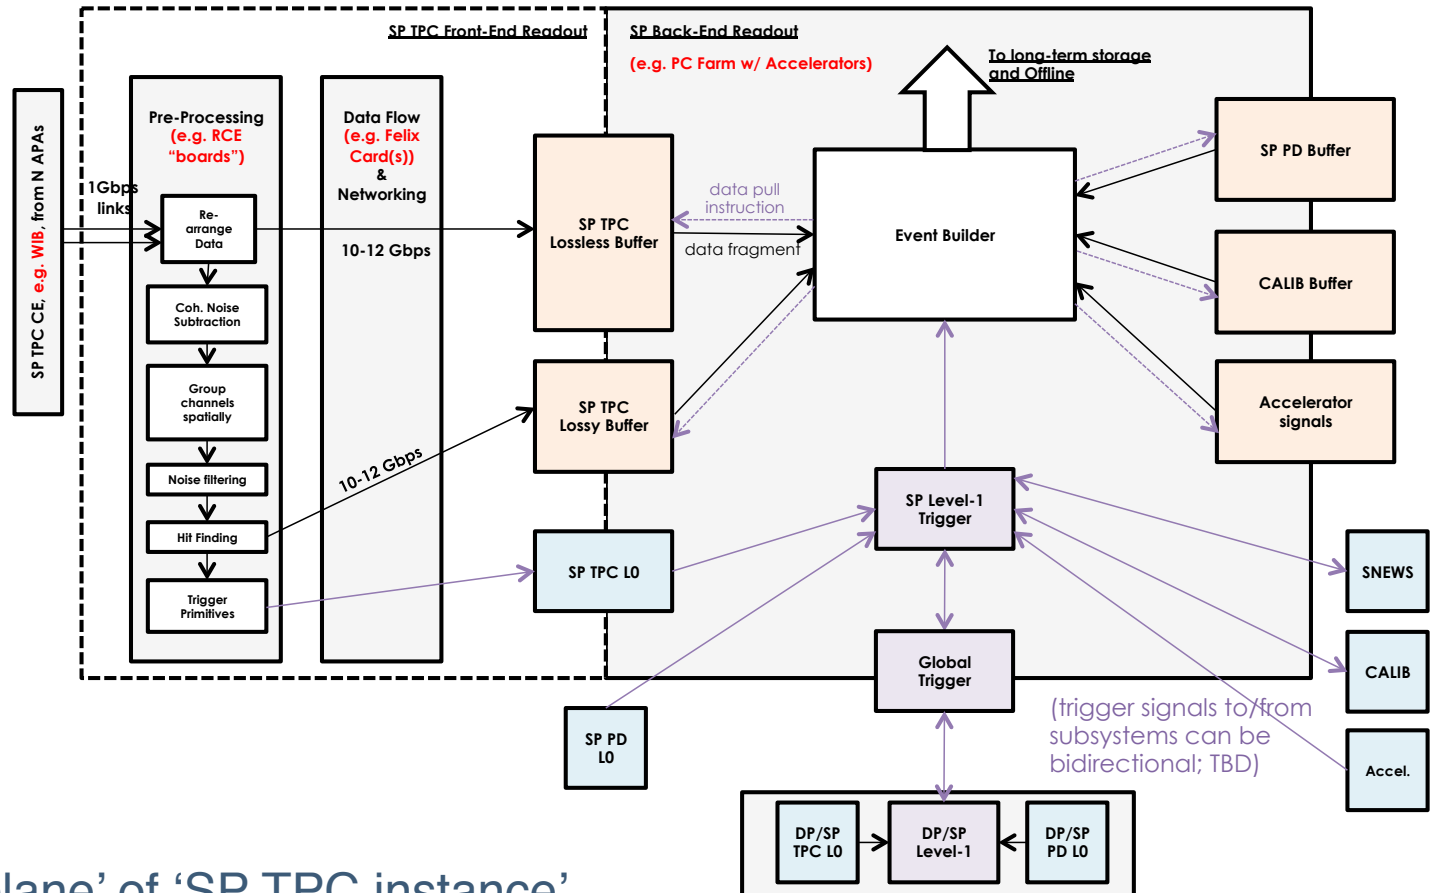
- Around 30 institutes in Europe, Japan, South America, US
 - Healthy level of participation: >50 active individuals, and growing weekly
- Divided into parallel working groups for the ante-TP era
 - Architecture / hardware & interfaces / computing / data selection / infrastructure and facilities
 - Plus ongoing DAQ simulation group, as our contact point to physics requirements
- Two productive workshops in our six months of existence
 - The division of long term interests / resources is becoming more clear
 - The Technical Proposal process has helped us to organise and find roles



Top-Level Requirements

- Single, scalable system for all detector modules
- Synchronisation to 10ns (within module) / 1us (between modules)
- Capable of buffering full raw data for 10s of seconds
- Zero deadtime under normal conditions (including SNB)
- Sufficient data reduction to achieve a maximum 30PB/year storage target
 - For four detector modules, asymptotically – but allow for this from day one
- Conform to physics-driven data selection requirements
 - >99% eff. for beam ν with $E_{\nu} > 100\text{MeV}$
 - >99% eff. for atmos. ν , nucleon decay, cosmic rays, with $E_{\nu} > 100\text{MeV}$
 - >90% eff. for galactic SNB, >90% eff. for triggered SNB ν with $E_{\nu} > 10\text{MeV}$
 - <10% of total data volume arising from SNB false trigger rate
- Conform to experiment infrastructure / integration requirements
- Many lower-level technical requirements; under review in lead-up to TDR

Dataflow Outline



- ‘Dataflow plane’ of ‘SP TPC instance’

- There are also control, trigger command and fast control planes
- Each plane partitionable at detector element level for commissioning, debugging and tests

Key Interfaces

- All interfaces now agreed, formally documented, subject to change control
 - In some cases well-defined technical options exist, final selection in TDR
- DAQ lives entirely in CUC, galvanically isolated from detectors
- Dual-phase electronics (TPC + photons)
 - Compressed raw data samples on commercial optical 40GBE / 100 GbE
- Single-phase electronics
 - Photons: similar arrangement to DP, ideally use similar protocol and format
 - TPC: uncompressed raw data, at 1Gb/s or 10Gb/s optical, custom protocol
- Timing and synchronisation (single timing domain across all detectors)
 - GNSS → CERN White Rabbit → PDTs (SP) or uTCA backplane (DP)
- Offline computing
 - Data handover point is WAN link to FNAL; ~300TB (3 day) buffer handles fluctuations, outages
- Conventional facilities, slow control, etc
 - Baseline requirements for underground and surface elements of DAQ now documented
 - Details will evolve as final hardware / computing design is finalised towards TDR

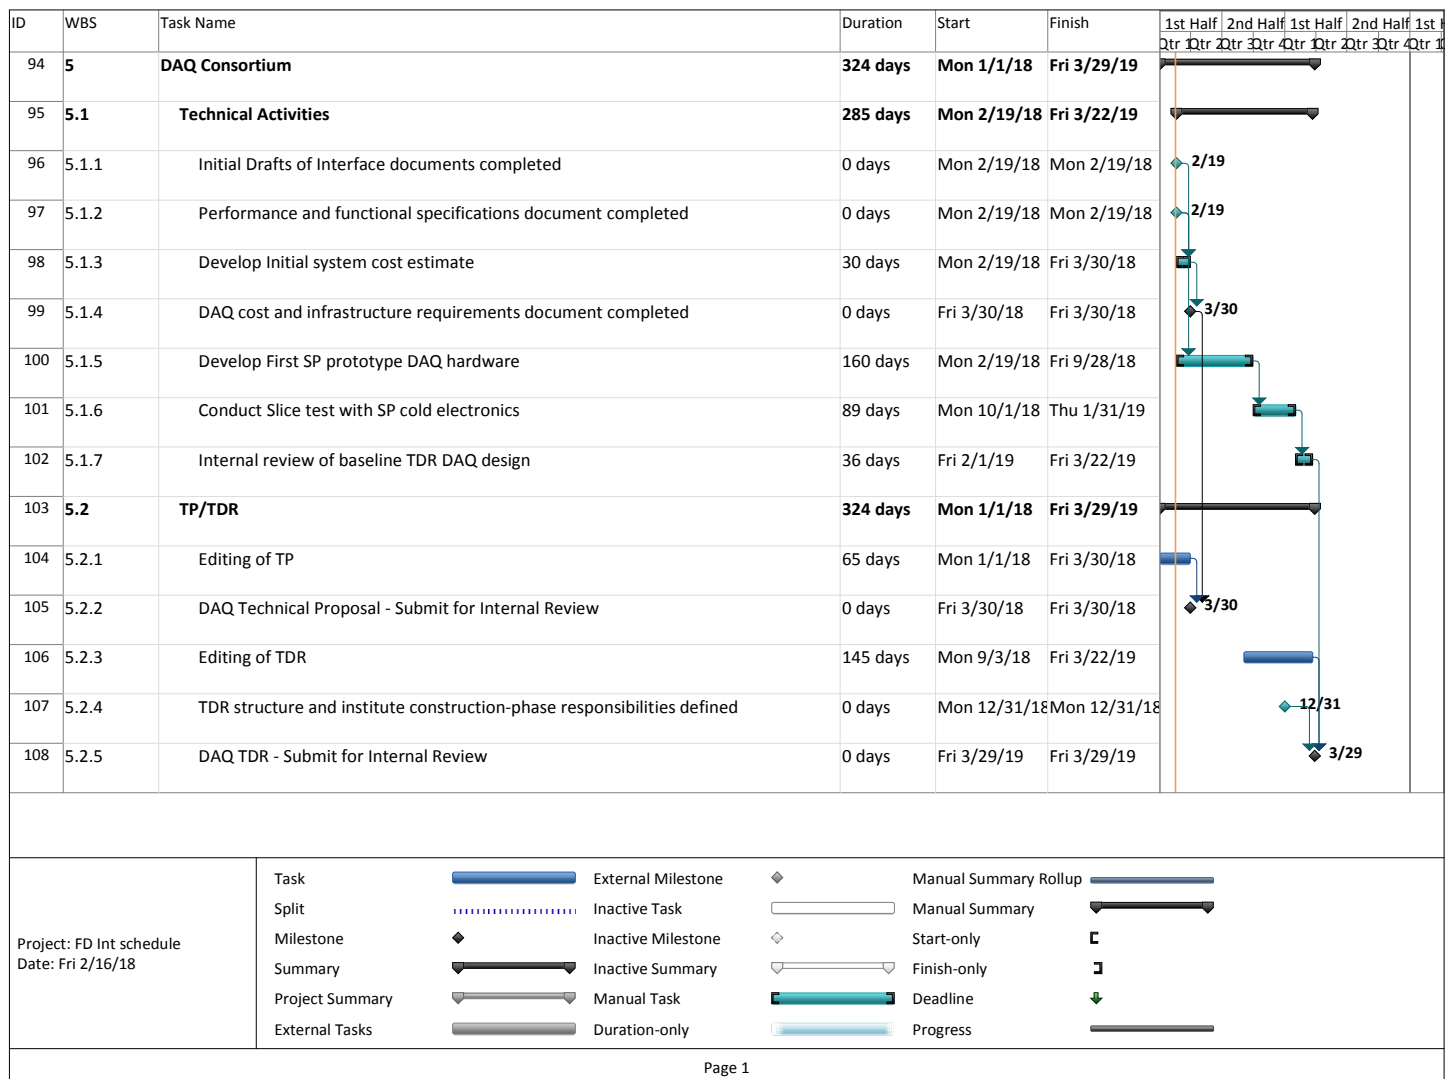
Agreed Design Principles

- A single, scalable system across all detectors
 - Allows use of common components, cross-triggering possibilities
- Capable of recording and storing full detector information
 - Invaluable for both detector studies / commissioning and physics
- Sized with significant conservatism for first module
 - Fixed infrastructure (e.g. power, network) sized for four modules from start
 - Allows a large safety factor in first year of running, when things are most uncertain
 - Keep all data selection decisions as simple as possible, at the expense of selectivity
 - As we learn about detector performance, will be able to be more selective
 - Where data reduction is needed, emphasise trigger cuts rather than zero suppression
 - Always keep the possibility to add data processing capacity if needed (e.g. L2 cluster)
- Criteria for design choices:
 - Robustness → Scalability → Ease of deployment and commissioning → Ease of design and construction → Operation costs → Capital costs
 - These will guide our choice of technical solutions for the TDR

Key Decisions before TDR

- Noise assumptions for sub-detectors
 - What are reasonable assumptions? What are the worst-case assumptions? ProtoDUNE...
- Final data rate to offline computing; DAQ parameters for SNB physics
- Physical location and arrangement of DAQ system
 - For now: split system at event builder, small surface computing, but full surface / remote ops
 - Limited data (<100Gb/s) up shaft, preserve possibility to send substantially more data
- Real-time filtering / compression / trigger primitive generation
 - How much functionality in FPGA firmware, how much in CPU / GPU ('software-centric system')?
- RAM buffer and non-volatile buffer implementation
 - >10s trigger latency buffer, at least this much NVM. Custom hardware, or commodity computing?
- Slice test strategy and location
 - Where are our bases of operations for test, integration and burn-in in 2019? In 2020? In 2021?
- We have 'working baseline models' for almost all of these questions
 - Fully expect to challenge these via further phase of R&D, leading to demo systems by early 2019
 - Final baseline will be decided pre-TDR, via collaboration review

Schedule



Pre-TDR Milestones

- M1 (Dec 2017) Interface documents complete (DONE)
- M2 (Jan 2018) Functional specifications complete (DONE)
- M3 (Mar 2018) Cost and infrastructure requirements complete
 - 90% done - not in TP, but ready for corresponding cost documents
- M4 (Mar 2018) **Technical Proposal**
- M5 (Aug 2018) Preliminary (internal) **review** of TDR baseline
- M6 (Oct 2018) First prototype HW / FW / SW available
- M7 (Nov 2018) TDR structure and institute responsibilities defined
- M8 (Jan 2019) Slice demonstrators complete
- M9 (Feb 2019) Full (external) **review** of TDR baseline
- M10 (Mar 2019) **Technical Design Report**

Risks and Concerns

1. Insufficient expert personnel to conduct and support project

- This is a large and complex system, need commitments from many experts
- Strategy: consolidate and expand consortium during R&D phase – we need more people

2. Detector noise specifications not met

- Strategy: provide horizontal (more processors) and vertical (more stages) scalability

3. Excess SNB trigger rate from instrumental effects

- Strategy: leave provision for L2 post-event builder processing
- Could include full or partial event reconstruction with induction planes

4. Calibration requirements grossly exceed offline data rate

- Strategy: near-line analysis for calibration in L2 cluster

5. Power / space requirements exceed CUC capacity

- Leave provision for all data to be shipped to surface
- Will require WDM on fibres – technically feasible, but expensive

6. Technical complexity results in lack of robustness

- Strategy: carefully defined demonstrator projects at increasingly scale

Response to Recommendations

- 2017-150: DUNE management should ... identify any additional need in the effort, tasks and expertise ... directed to the ProtoDUNE-SP DAQ group.
 - ProtoDUNE-SP DAQ technically outside DAQ consortium
 - BUT: Convergence of ProtoDUNE and DUNE FD DAQ clearly essential in 2018
 - New DAQ collaborators are strongly motivated to learn from ProtoDUNE
 - This is key goal of the demonstrator programme, centred at CERN
- 2017-151: Lessons learned from ... protoDUNE should be formally documented in real time...
 - Large overlap between active ProtoDUNE and DUNE FD DAQ personnel
 - Need for (living) documentation reinforced with ProtoDUNE leadership
- Demonstrator programme will be closely coupled to, and a seamless transition from, ProtoDUNE
 - i.e. continuation of ProtoDUNE facilities beyond 2018 essential part of DAQ project

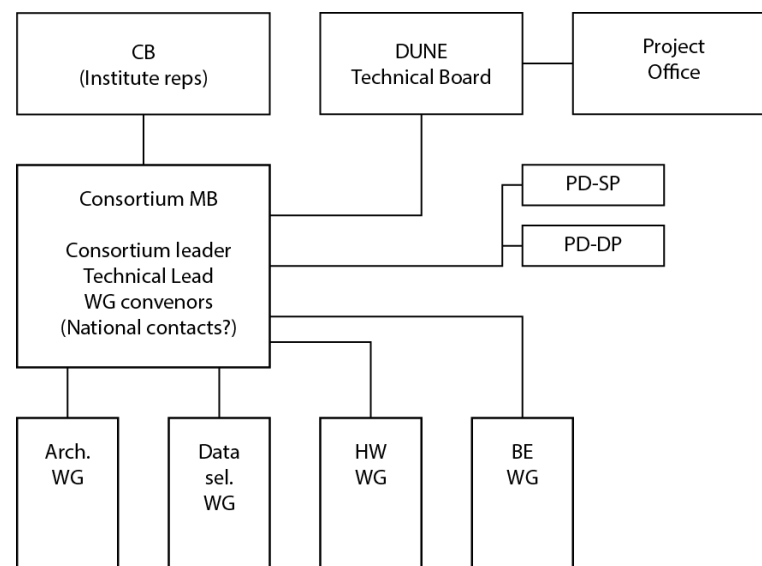
Summary

- In the last six months:
 - Inaugurated a strong and growing DAQ consortium
 - Developed a baseline DAQ architecture for all modules
 - Identified potential technical solutions to be carried to next R&D stage
 - Including a fully documented design, based on ProtoDUNE, also basis for costs
 - Established draft WBS & schedule, began assigning responsibilities
- What's ahead of us:
 - Deliver the Technical Proposal – now in end game
 - Agree R&D / demonstrator plan for the pre-TDR phase – achieved
 - Convergence and overlap with ProtoDUNE during 2018 is a project requirement
 - Make key decisions, based on clear process and criteria
 - Deliver Technical Design Report with a full construction plan
 - Continue to build world-class team around a challenging project

Backup

Consortium Structure

	Institution	Contact
France	Institut de Physique Nucleaire de Lyon (IPNL)	Dario Autiero
Japan	National Institute of Technology Kure College	Seiji Kasai
Japan	Iwate University	Shinya Narita
Japan	KEK	Takuya Hasegawa
Netherlands	Nikhef	Paul de Jong
CERN	CERN	Giovanna Lehmann Miotto
United Kingdom	University of Birmingham	Alan Watson
United Kingdom	STFC Rutherford Appleton Laboratory	Alfons Weber
United Kingdom	Univ. of Warwick	Gary Barker
United Kingdom	University of Liverpool	Karol Hennessey
United Kingdom	University of Oxford	Giles Barr
United Kingdom	University of Sussex	Simon Peeters
United Kingdom	University of Bristol	David Newbold
United Kingdom	University College London	Ryan Nichol
United Kingdom	Edinburgh University	Franz Muheim
USA	Brookhaven National Lab	Brett Viren
USA	Columbia University	Georgia Karagiorgi
USA	Duke University	Kate Scholberg
USA	Fermi National Accelerator Lab	Kurt Biery
USA	Iowa State University	Amanda Weinstein
USA	University of California (Davis)	Bob Svoboda
USA	University of California (Irvine)	Michael Smy
USA	University of Minnesota (Duluth)	Alec Habig
USA	Notre Dame University	John LoSecco
USA	Pacific Northwest National Lab	Eric Church
USA	University of Pennsylvania	Josh Klein
USA	SLAC National Acceleratory Laboratory	Mark Convery
USA	South Dakota School of Mines and Technology	Juergen Reichenbacher



Costs Snapshot

- These numbers are not public, and explicitly subject to change

	SP	DP	Common	Total
Fibres	740.9237	132.8	0	873.7237
Hardware	4064	0	0	4064
FE computing	978.5	710	0	1688.5
BE computing	203	116	2	321
Service computing	45	45	6	96
Timing	95.78	0	24.4	120.18
LAN	82	106	5	193
Facilities	105.6	105.6	245.4	456.6
	6314.8037	1215.4	282.8	7813.0037

Data Selection Context

- Dominant sources of `signal':
 - ▶ Radiologicals (^{39}Ar , ^{42}Ar , ^{208}Tl , ^{210}Tl , ^{224}Rn ...)
 - ▶ Cosmic Rays (~4500/day in 10 ktonne)
 - ▶ Front-end calibrations
 - ▶ Radioactive source calibrations
 - ▶ Laser source calibrations
 - ▶ Detector noise and other anomalies
- Operating principles:
 - ▶ Minimise potential bias to data
 - Want to accept anything that is noticeably different from noise or radio-accidentals
 - ▶ Be as simple as possible
 - ▶ Be as flexible as possible

Data Selection Concept: SP TPC

- Form decisions hierarchically
 - ▶ **Local** view (per APA?), collection wires only – extract summary data
 - ▶ **Module** view, i.e. of whole SP TPC – find statistical signatures
 - ▶ **Global** view: bring together information from SP TPC / PD, DP TPC / PD, external
 - This is where the ‘event’ (i.e. contiguous range of data in space / time)is defined
- Data ranges – what is the *simplest thing we can do?*
 - ▶ For most events: keep entire detector for ~2 drift times
 - ▶ For SNB, stream entire detector to NVM for ‘extended period’
 - Implies memory system with sufficient bandwidth to keep up
 - ▶ Keep primitives continuously for sub-threshold sensitivity
- Much more sophistication is clearly possible, and in the end necessary
 - ▶ We will bring this into play as we learn more about our detector
- We now need a corresponding strategy for other sub-detectors

Technical Proposal: SP

22	7 Data Acquisition System	39
23	7.1 Data Acquisition System (DAQ) Overview (Georgia Karagiorgi)	39
24	7.1.1 Introduction	39
25	7.1.2 Design Considerations	39
26	7.1.3 Scope	41
27	7.2 DAQ Design	42
28	7.2.1 Overview (Giles Barr)	42
29	7.2.2 Local Readout & Buffering (Giles Barr & Brett Viren)	42
30	7.2.3 Local Trigger Primitive Generation (Josh Klein & Brett Viren)	42
31	7.2.4 Dataflow (Giles Barr & Josh Klein & Brett Viren)	44
32	7.2.5 Data Selection Algorithms (Josh Klein & Brett Viren)	44
33	7.2.6 Timing & Synchronization	45
34	7.2.7 Computing & Network Infrastructure (Kurt Biery?)	45
35	7.2.8 Run Control & Monitoring	45
36	7.3 Interfaces (David Cussans & Matt Graham)	45
37	7.3.1 TPC Electronics	45
38	7.3.2 PD Electronics	45
39	7.3.3 Offline Computing (Kurt Biery?)	46
40	7.3.4 Slow Control	46
41	7.3.5 External Systems	46
42	7.4 Production and Assembly	46
43	7.5 Installation, Integration and Commissioning (David Newbold & Alec Habig)	46
44	7.5.1 Installation	46
45	7.5.2 Integration with TPC/PD Electronics	46
46	7.5.3 Commissioning	46
	7.6 Safety (David Newbold & Alec Habig)	46
	7.7 Organization and Management (David Newbold & Georgia Karagiorgi)	47
	7.7.1 DAQ Consortium Organization	47
	7.7.2 Planning Assumptions	47
	7.7.3 WBS and Responsibilities	47
	7.7.4 High-level Cost and Schedule	47

Editors: J. Brooke, B. Viren

Data Volumes

Source	Annual Data Volume/ 10 kt	Assumptions
Beam interactions	20 TB	700 vs+700 dirt μ s; everything read out for 5.4 ms (no ZS); 10 MeV threshold in coincidence with beam time; include cosmics
Cosmics (+atmospherics) Scheme 1	10 PB	All wires in 5.4 ms window around HE event
Radiologicals	< 1 PB	Weighting scheme for SN bursts and other LE events; dump all 10 s for each burst trigger; tune fake rate to be < 100/year.
Front-End cals	200 TB	Worst case of measuring every single ADC bin with 100 measurements/point; four times/year
Radioactive source cals	100 TB	Source rate < 10 Hz; only one APA readout; PDS is negligible; full readout window per tag; no ZS
Laser cals	200 TB	1x10 ⁶ total laser pulses; tight ZS for both induction and collection; ½ of all wires in TPC illuminated
Random Triggers	60 TB	Same as cosmics scheme; rate is 45/day
Trigger Primitives	< 2 PB	Only collection wires; 12 b/primitive; 4 primitive types; ³⁹ Ar dominates;

- 30PB/year (asymptotically) plausible – of equal importance, we have knobs to adjust
- ‘Conservative’ sizing allows >10PB/year in year #1 as needed
- Refinement of estimates an active process, continuous interaction with physics

SNB Fake Rate vs Coverage

